"Risk Measurement and Management: A Financial Engineering Roundtable."

On March 6, 1998, the Wharton Financial Institutions Center held a Financial Engineering Roundtable on Risk Measurement and Management in financial institutions. The event, organized by Anthony Santomero and Francis Diebold, was supported by Oliver, Wyman & Company and brought together leading academics, representatives of major trading houses, and several well-known consultants in the area to discuss the limits of existing systems and the need for advancements in this field.

The event centered on the "state of the art" or science available to major institutions that have committed their organizations and substantial resources to a large-scale effort in trading activity. It began where most conferences end in that the focus was not the operational details of current risk management systems; rather, the conference's objective was to explore new frontiers and to question existing methodologies in light of available empirical evidence.

The format was consistent with the event's title. There was little pure presentation and more open discussion than is generally the case at such events. Over the course of the day, the Roundtable covered a total of five sessions: the challenges of measuring and modeling the tails of asset return distributions; horizon problems in risk management; mechanisms for adequate risk adjusted performance measurement and compensation; appropriate methodologies to measure the market's estimate of asset risk; and the challenges of pricing and measuring the risk of the newer assets surfacing in the capital market every day. For each session, several leading researchers were asked to outline the issues to be discussed and the state of industry knowledge, and then the floor was opened for discussion of what this work means for risk management systems and processes.

The event was open by invitation to senior risk managers from the major trading institutions of the industry. Representatives from such firms as JPMorgan, Goldman Sachs, Bear Stearns, UBS, Deutsche Bank, CIBC, Morgan Stanley, Toronto Dominion, Royal Bank, BlackRock, Santander and First Union were present at the roundtable.

What follows is a brief "snapshot" of these discussions. We welcome your questions and comments. Please contact Anthony Santomero, the Center's Director, or any member of the Financial Institutions Center team with your questions and thoughts.

Measuring and Modeling the Tails of Distributions

Discussion Leaders:
-Paul Embrechts, ETH (Swiss Federal Institute of Technology) Zurich
-Casper G. deVries, Erasmus University
-J.Huston McCulloch, Ohio State University

Risk management is concerned with estimating tail probabilities and quantiles of profit-loss distributions, especially at the extremes. There is a need for methods for estimating conditional probabilities concerning tail events and answering the following question: given that a loss beyond Value-at-Risk (VaR) is incurred, how great do we expect the loss to be?

Paul Embrechts opened the session with the view that Extreme Value Theory is well suited to the problem of accurate distribution measurement, if interest centers on the tails of such distributions. This approach has been well developed in other areas and could also be employed in financial risk management measurement issues.

Given that it is known that profit and loss data are heavy-tailed, skewed, and estimate rare events, how can one accurately measure VaR? And then having found this number, can anything else be said about this exposure that is relevant for the management of risk?

Extreme Value Theory gives us a way to estimate high quantiles of risk. Given distributions that are fat-tailed, it is effective at fitting a distribution far out in the tails. The approach then calculates high quantile risk, which can be translated to a VaR measure. Beyond this, Extreme Value Theory allows for the estimation of the expected value of a loss above the VaR measure.

He argued that VaR by itself is not a management tool in the conventional sense: it can be inconsistent if the underlying distribution is different from multivariate normal. Also, since the expected value of the supremum of loss is being sought, VaR is by itself problematic.

The general opinion of the discussion was that stress scenarios are better than VaR and that looking at the expected loss, given a loss event beyond a certain level, gets closer to addressing true management concerns. Management is not just looking at a frequency and a loss quantity, but needs to know the expected size of these losses. Extreme Value Theory is useful when interest is centered on extremal events, and works better than other empirical estimation procedures and guesswork.

However, there are limitations to this method. Calculating extreme values becomes difficult for portfolios of large dimensions. Also, computational sources are not automatic, and thought must be put into the building of the model. Finally, one cannot make predictions using Extreme Value Theory, one can only make assessments based on the observed data.

The second speaker, Casper G. deVries, discussed the calculation of Value-at-Risk. Current VaR methodology can be classified into two types:

- Parametric modeling of conditional volatilities (such as J.P. Morgan's RiskMetrics); and,
- Non-parametric prediction of unconditional volatilities, such as techniques based on historical simulation or stress-testing methods.

Historical simulation cannot go beyond the empirical sample data, while the variance of stress tests is dramatically large. And, while the use of a normal distribution assumption can go beyond the sample, the data are distinctly non-normal. For example, the normal misses the peakedness of the data in the middle and misses the outliers in the tails.

What can be done to deal with the heavy tails evident in the data? It is known that as the sampling of a random variable becomes large, the tails decline. With a normal distribution they decrease exponentially, and hence have short tails. With Extreme Value Theory they decrease by a power. This is a crucial difference, and explains why the normal model fits badly at the outer parts of the tail.

In historical simulation, the empirical distribution has little information about the tails because there are so few observations in the tails. The normal curve declines exponentially at the tails and misses out on the tail features. The power approximation, however, nicely captures in-sample properties from heavy tails, while still being able to go out-of-sample. This result advocates the semi-parametric method: no information is necessary about the middle, just the tails. Tails are modeled parametrically, but hyperbolically for only the largest outcomes, not by exponential type approximation. For the middle the empirical density can be used, which is superior because there are many observations in this area.

deVries noted several caveats to the implementation of the extreme value method. In applications that are concerned with extreme outcomes, as is typical for VaR analysis, the data set has to include a sufficient number of extreme events in order to obtain an accurate prediction of VaR. It is, however, not that difficult to implement the tail estimation procedure.

Nevertheless, the current set of Basle requirements still provides disincentives for more reliable VaR models. This stems from the regulators' tendency towards more conservatism in their estimation of risk. deVries expressed his belief, however, that considerable improvement of current VaR models is possible by means of techniques that focus explicitly on the properties of extreme return fluctuations.

J.Huston McCulloch led the final discussion on the subject of the Financial Applications of Stable Distributions.

Financial asset returns are the cumulative result of many investors' decisions. The Central Limit Theorem tells us that if the sum of a large number of independent, identically distributed random variables has a limiting distribution, that limiting distribution must be a member of the stable class, of which the Gaussian, or normal, is the most familiar. The fact that the Gaussian is the most familiar and tractable of these stable distributions has caused it to be used frequently as a model for asset returns. However, returns have much "fatter tails" than is consistent with normality, which leads one to consider the possibility that non-Gaussian stable distributions might be better models of financial returns.

If asset returns are governed by the infinite-variance stable distributions, there is a greater possibility of the occurrence of extreme events: life is much riskier than is implied by the Gaussian model. However, the CAPM works as well in the infinite variance stable cases as it does in the normal case, and the Black-Scholes formula may be extended to the non-Gaussian stable cases by means of a utility maximization argument.

From the three discussions in this session, it was concluded that stress testing may be a preferable method for assessing risk from a banking point of view, but does not allow the prediction of changes at a large, previously unseen level. Value-at-Risk focuses on the frequency of events, but not on their severity. The advantage of methods that allow for fat-tailed distributions is that they allow a more accurate view of what happens in cases that are outside of empirical observation. Also, it was pointed out that there are some practical problems with extreme value theory: multivariate cases are difficult to manage, and are as yet un-automated, but these problems can be solved.

Horizon Problems

Discussion Leaders:
-Francis X. Diebold, University of Pennsylvania
-Torben G. Andersen, Northwestern University
-David A. Hsieh, Duke University

Frank Diebold opened the discussion of this session by highlighting the pitfalls of scaling as a method for converting short-horizon volatility to long-horizon volatility. Short-horizon risk measures, such as daily or intraday measures, are commonly converted to other horizons by means of scaling. For example, to obtain 10-day volatility from one-day volatility, the common approach would be to multiply one-day volatility by the square root of 10. But Diebold warned that scaling mistakenly assumes that high-frequency returns are independent, identically distributed (iid) random variables.

Doing the aggregation more formally is a complex task, and aggregation formulae are only available for limited types of models. It would be simpler if scaling, in spite of its not being strictly appropriate, nevertheless gave an accurate approximation. Unfortunately, scaling by definition magnifies volatility fluctuations, whereas theory suggests that aggregation should reduce volatility fluctuations.

One way to get around this problem is to use a non-parametric scheme for getting a sense of the forecastability of volatility, i.e., how quickly and with what pattern does forecastability decay as the horizon is lengthened? This method is not dependent on fitting a specific model, but rather, comes out of evaluating confidence intervals. The results indicate high volatility predictability at very short horizons, but the predictability decays quickly, and there are substantial differences in how soon volatility predictability falls off for different products.

Next, Torben Andersen spoke on the topic of high frequency data and volatility processes. Looking at high frequency data for volatility processes brings up more complications for the risk manager. Standard approaches do not work, as stochastic volatility models fall flat for intraday data. While day-to-day shifts are predictable in one to five day horizons, intraday modeling suffers from the large effect calendar events have on the data. For example, jumps of 25% are very rare on an inter-day level, but jumps of 25% and higher are not uncommon intraday, following an announcement or near the open/close of a trading day. The risk manager must be able to discern, given a large shock, whether it is simply a reaction to new information, a change in volatility, or just a random shock.

Finally, David Hsieh discussed horizon problems and management styles. Here, he asked the provocative question: What is the time horizon over which one should be considering issues of risk management? Standard Value-at-Risk measurements take the current position and give an estimate of how that value will change over time. The implicit (and misleading) assumptions are that the position's value change is relevant over the time interval of estimation and that it remains constant over the period.

When determining the frequency of appropriate measurement of Value-at-Risk, it is important to take into consideration the investment style and objectives of the portfolio managers, including their trading strategies and their use of leverage. In other words, one would need less frequent updates on the position of a portfolio managed by someone with a buy-and-hold strategy and more frequent updates about a fund whose manager has a more dynamic strategy.

Spurred by these presentations, the Roundtable discussion concluded with the following observations:

- Practitioners use end-of-day measurements to calculate risk even though they know that positions can be quite different the following morning. The tendency is to put in place policies that govern traders' activities rather than policies that govern actual positions. In practice, the observations of horizon problems were seen as useful for solving specific problems, but difficult to address in a generalized way.

- It is critical to be aware of the underlying data used in any risk management system. For example, even if volatility is not moving around, returns may be: an estimation made from a moving average might result in one that shows volatility where none exists. Further, the time of day may have a huge effect on the results of the analysis.

- It is also important to keep in mind that total risk includes the volatility of many asset positions, and correlations across markets are essential. However, the caveats concerning data are equally applicable here.

- Finally, it is critical to be aware of the trading styles of the managers making the decisions, and to take into account how fast capital can be moved. It is best to have a planned strategy in a given scenario.

Risk Adjusted Performance Evaluation and Compensation

Discussion Leaders:
-Armen A.Avenessians, Goldman Sachs
-Till M. Guldimann, Infinity Financial Technology, Inc.
-Eli Wachtel, Bear Stearns & Co.

In this session, the panelists considered various solutions to the questions: How does one monitor trader performance in a way that is adjusted for risk? Are there feasible compensation schemes that encourage the trader to optimize the risk-return trade-off?

It is difficult to give an adequate assessment of return, as it is the result of the combined efforts of traders, salesmen, and quantitative pricing agents. One can look at the compensation issue as a market for traders. Traders tend to know their market value, but how is it determined for the firm?

Few firms have focused on this market in terms of a return-versus-risk compensation scheme; most use compensation schemes that are return-based, and impose risk controls instead. However, any control system that the firm imposes, such as Value at Risk, will be subject to gaming, as traders seek to optimize the benchmark against which they are compared.

Objective measures of risk are important but difficult to implement at the operational level. In practice, risk assessment can frequently be ad hoc and difficult to defend. The firm's assumptions about risk can often go against the trader's positions. It would be difficult to determine compensation numbers based on these assumptions.

Moving from an environment of simply controlling risk to one of managing risk presents some unique challenges. In many of the major trading organizations, risk is taken at the lower part of the organizational pyramid, while controls must be imposed from the top. Yet, it is necessary to find the means to manage risk-taking in a way that optimizes return on risk.

The current management model imposes the level of risk that an individual is allowed to take, and then ensures that he or she remains within that limit. Then, it evaluates the returns achieved in the context of the risks taken.

Limits can be expressed in terms of Value at Risk, but it is more useful to report P&L and VaR together. Furthermore, the firm should provide these data for each level of the organization, and be able to put these numbers in an historical perspective.

From the point of view of the trading manager, intuitive measures of trader capability are at least as important, if not more important, as objective measures of risk. VaR provides one good measure of risk in speculative positions, but in most cases the manager must go beyond VaR to gain sufficient insight about the trade. The manager must have a thorough understanding of the details of the trade, assessing risk based on expected profitability and expected volatility. For this, the manager must bring a significant amount of experience to the job. And, of utmost importance, the trading manager must make sure that the trading desk never makes a bet it cannot afford to lose.

Looking at the Sharp Ratio is useful to a certain extent when the objective is to maximize revenue relative to risk, steady and consistent profits are more valuable for the manager than large but erratic ones. The objective is to maximize return with minimum volatility.

In conclusion, it was indicated that the role of VaR is tied less to trader compensation, and more to the policing of risk. The regulatory push has been more towards using VaR to contain institutional loss, rather than in the allocation of resources internally. However, VaR is most useful at the management level, allowing for a better understanding of risk and greater transparency between different activities of the firm.

Using Options Prices to Assess Market Risk

Discussion Leaders:
-Andrew Lo, M.I.T.
-David Bates, University of Iowa
-Jose Campa, New York University

Andrew Lo opened this session with a discussion of investor perception and choice in risk selection. In a series of well-known experiments on loss aversion, it is known that when faced with a decision that carries some risk, people will consistently choose to minimize their downside outcome. This suggests that consumers may not always make economically rational decisions in a risky environment, and one should not expect observed prices to be consistent across markets or over time. It is therefore difficult to extract preferences from observed market prices of various financial assets.

Nonetheless, it is important to understand how these uncertainty preferences can be extracted from prices of securities. To do so, one can not simply extract volatility measures from the data, but one must go from estimates obtained from the data to the state price densities.

Can one use the empirical data from options prices to assess market risk? Is information reliable for measure of VaR or risk management in general? These questions were addressed by David Bates.

It is recognized that the Black-Scholes model is an inaccurate way of extracting estimates of risk and implied volatility. It returns a different implied volatility for different strike-prices and different maturities, whereas if the model were accurate we would get the same implied volatility.

This is true for all the different methods that are used to extract these implicit distributions from option prices. Examples include models driven by stochastic processes, numerical derivatives, constrained cubic splines, and entropy-based approaches. Looking only at risk distribution at optimum maturity, the choice of methods used does not matter much. They all fit both the option process and the second derivative reasonably well. Furthermore, if interest is only in volatility, then Black-Scholes is about as good as any other model. The results are not unbiased, but the bias is not very significant.

However, it does matter which model is chosen when interpolating. When volatilities are not flat, interpolating across maturities will result in large interpolation errors. At higher moments, extrapolation is taken even further, and become even more prone to error.

Nonetheless, Bates argues that option prices do contain useful information about market risk, enough so that it is worthwhile to look at this information when performing risk assessment. However, this information cannot be taken at face value. After this data is processed, it must be cross checked against the time series, in order to decide what weight to give the information relative to other information, such as GARCH and other time series assessments of risk that are in place.

Jose Campos continued this line of discussion by noting that the methodology for assessing market risk using option prices is well established. All methodologies use some form of the idea that the second derivative of call prices is proportional to the risk neutral density. Implementation of this methodology can be done using the volatility quotes market to get the interpolated volatility function, and then using the Black Scholes option formula to take the second derivative.

A second approach is distribution-based, taking a distribution with certain parameters, for example the implied binomial distribution or a mixture of log-normals. Then they are parametrizied to fit observed option prices plus the forward rate or other information.

Still other methods include the implied binomial tree, both untrimmed and trimmed to remove negative probability, and log-normals.

There is skewness in the results obtained from calculating Black-Scholes "at the money". If only the distribution of the first moments are needed, then all models do very well relative to Black–Scholes, and there are not many differences between the alternate methods. If, however, higher moments are needed, skewness and kurtosis become problematic.

At higher moments all of the methods yield very different results. Moments are very sensitive to the way they are estimated, and it is difficult to tell which case is the correct one – we do not know the true kurtosis.

It was concluded that all methods seem to provide better fits to data than Black-Scholes. Higher moments are, however, poorly estimated because they are extrapolations beyond the first moments. Despite this, useful information can be gained from these distributions. They are useful in computing the maximum probability of certain realizations. One can look at the measures of the probability of remaining within a band. For example, the shapes of the implied foreign exchange distributions change dramatically over time and one can use the distribution to manage an international portfolio better than by simply looking at options or at different stocks.

New Directions

Discussion Leaders:
- Neil Doherty, University of Pennsylvania
-Richard Lyons, University of California, Berkeley,
-Gary Gorton, University of Pennsylvania

The program ended with a session on new markets and instruments. The questions here were both how are these markets evolving and the pricing methodologies being applied, as well as how should the risk manager think of these instruments fitting into the overall risk control mandate.

Catastrophic Risk Instruments

Neil Doherty opened with a discussion of the various types of CAT bonds emerging in the market. He noted that recent catastrophic events, such as Hurricane Andrew and the Northridge earthquake have imposed costs on the insurance industry of an order of magnitude not thought possible only a decade ago. As a result, catastrophe risk instruments have emerged as a new asset class, as insurers seek ways to diversify the risk through capital market risk transfer.

Instrument types vary from straight debt with interest payments that are dependent upon an index of loss associated with catastrophic events, such as hurricanes or earthquakes, to equity puts, where certain events trigger the conversion of the debt instrument to equity. The former may be principal protected or not, depending on instrument design or the tranche purchased. The latter debt instrument, on the other hand, has the characteristic that it can be converted to equity after a defined event, but the option is held by the issuing firm rather than the bondholder. This provides a partial hedge.

This set of instruments performs differently from one another and from standard financial assets in many dimensions. However, some important asset-specific features warrant particular attention. Some have substantial credit risk, while some do not. A concern for counterparty credit risk is understandable, as the events that they are hedging against with these instruments are the same events that would put stress on the counterparties.

A second dimension is the determinant of a claim. Two versions exist, viz., portfolio specific coverage or index based claims. These in turn cause two different types of concern. With explicit firm level coverage, the investor has a moral hazard problem. The insurance companies' portfolio of coverage determines the probability distribution of loss associated with some of these instruments but the financial investor picks up the bill.

The alternative structure causes a basis risk problem. Here, the issuers' coverage is not directly related to firm specific losses but coverage is related to the movement of an aggregate loss index of some kind. Basis risk and moral hazard tend to counteract each other, and cause offsetting pressures on instrument design.

Essentially, the debate is whether you hedge against what the insured loses, i.e., an indemnity hedge, which is expensive because of the moral hazard problem, or against an index, for example, by aggregating insurance companies within a region, and taking an index of insurance company experience in that state as a trigger. Finally, there can be a parameterized trigger, which defines payoff as related to a parameter that is not open to manipulation, such as the intensity of an earthquake, the force of a hurricane, and so on. The underlying idea is to control the moral hazard problem.

Whatever the choice, the pricing of these instruments comes from an analysis of state contingent payoffs. These are obtained from one of several major modeling firms that calculate expected payoffs and the associated uncertainty. Their current attractiveness comes from the perception that the spreads on recent deals have been rich and the presumption that they have close to a zero correlation with other financial instruments.

Foreign Exchange Risk

Richard Lyons next reviewed the state of foreign exchange trading, emphasizing current trends in customer use, rather than new instruments and techniques. He argued that the current state of the market was the result of six important factors affecting firm level activity in foreign exchange risk markets:

1. Continuing and increasing emphasis on sources of value creation at the corporate level. Firms are, therefore, rolling back the decision process, and evaluating whether hedging is right for them, given their level of risk, leverage, and the transaction costs involved in hedging.

2. Shifting from risk minimization to profit maximization. There are two main categories of risk in foreign exchange, viz., current unsettled transaction exposure, and operating exposure, i.e., exposure to future net revenues from changing foreign exchange rates. Many firms see transaction exposure as easy to control, and operating exposure as difficult, and concentrate on the former at the expense of the latter. To manage the latter type of risk the firm must look at its pricing policy. In addition, he points out, the two might be correlated, requiring that the firm to evaluate and manage both simultaneously.

3. Increasing emphasis on performance evaluation and controls in this area as firms and their boards move towards greater policy articulation. Boards are increasingly asking their managers to state objectives for risk management activity and have their actions measured relative to successfully meeting these goals.

4. Increasing emphasis on using foreign exchange instruments in emerging markets to reduce risk from exchange revaluation and interest rate movements. Firms are looking at devaluation as a regime risk,

e.g., moving from fixed to flexible exchange rates.

5. Firms seem to be engaging in foreign exchange transactions increasingly because accounting rules are pushing them toward instruments that may be less appropriate for them. The argument is that firms might be heading towards the use of options when options may not be the right tool, precisely because they are being afforded hedge accounting. As these rules are sorted out, this issue may prove to be important.

6. Finally, there has been general ambiguity on the part of firms using options, as to whether they were hedging or risk taking activities. In general, the expectation is that more firms will decrease rather than increase their activities in this area. Many firms are finding that these activities are not adding value and that this hedged risk is not as significant as they had initially thought.

Credit Derivatives

The group next turned its attention to a presentation on credit derivatives led by Gary Gorton. He began by noting that the credit derivatives market has grown in interesting ways over the past several years.

One basic instrument is the credit default swap. In this case, two parties, Bank B, a protection buyer, and Bank A, a protection seller, sign a contract that says Bank B will make a payment to Bank A over some period of time as long as a given underlying security does not have a credit default. Suppose the bond then has a credit event. Then, and only then, will Bank A make a payment to Bank B. This payment could either be a fixed sum or, in the more likely case, it would be the face value of the bond minus the price of the bond for some period after the credit event. So, the security payment is triggered not as a function of price movement but as a result of a specific event.

Note that Bank B does not need to own the underlying security; the payoff is simply linked to the security because Bank B is interested in buying credit protection against the firm that issues the underlying bond.

There are a variety of different structures for these credit derivatives, of which the above is only one example.

The unsettled issues in this market generally related to the early stage of its development. For example, there are issues that need to be resolved about the contract. No standard contract currently exists and there are disputes about what the contract should say. The ISDA has proposed a contract that appears to have made things worse, not better.

The biggest disputes concern the definition of the credit event. Must it require default? For a simple example, there could be a repudiation of responsibility for Korean Development Bank bonds, even though the Korean Development Bank does not default. There are a long list of such potential cases that need to be negotiated in figuring out what constitutes the credit event.

Another issue centers around delivery. It may seem obvious that parties wish to cash settle contracts at maturity. However, in actuality many parties prefer physical delivery, a factor that can be a serious problem. For example, if a credit default swap is written and the company defaults, it may be difficult to obtain physical delivery. In an extreme case, the price of an instrument of a firm in Chapter 11 may diverge from the value of the underlying cash flow because of a need for physical delivery.

From a pricing perspective, these instruments behave quite differently from the underlying security. In fact, there are important issues in terms of pricing these contracts. For each swap, the cumulative probability of a default event must be estimated. And the default swap-price of the instrument after the event or is it an idiosyncratic event? Addressing these issues is a data-intensive operation.

The market has so far developed using the noisy information surrounding the use of bond ratings. This sidesteps many complicated issues associated with valuation using fundamental analysis, including the use of date and standard term structures are also random variables. Hence, the default probabilities, bond price, and interest rates are random variables that must be estimated simultaneously and consideration must be given to the fact that they may well be correlated. Is default risk a systematic event inaccurate data. One way currently employed to attack this problem is to use historical data, or to attempt to extract information from actual prices, but there are large problems inherent in these strategies. Ratings have not performed that badly in comparison and continue to be used until a better equally robust approach emerges.

To view further Financial Institutions Center research in this area you can access the following working papers at our web site, at http://fic.wharton.upenn.edu/fic/:

98-10	Pitfalls and Opportunities in the Use of Extreme Value Theory in Risk Management Francis X. Diebold, Til Schuermann and John D. Stroughair, April 1998
98-12	Financial Innovation in the Management of Catastrophe Risk Neil A. Doherty
98-16	Horizon Problems and Extreme Events in Financial Risk Management Peter F. Christoffersen, Francis X. Diebold and Til Schuermann
97-34	Converting 1-Day Volatility to h-Day Volatitlity: Scaling by Root-h is Worse Than You Think Francis X. Diebold, Andrew Hickman, Atsushi Inoue and Til Schuermann, July 1997
97-37	Evaluating Density Forecasts Francis X. Diebold, Todd A. Gunther and Anthony S. Tay, August 1997 forthcoming, International Economic Review
97-45	How Relevant is Volatility Forecasting for Financial Risk Management? Peter F. Christoffersen and Francis X. Diebold, October 1997
96-25	An Integrated Approach to Hedging and Pricing Eurodollar Derivatives Robert Jarrow and Stuart Turnbull, May 1996
96-40	Bank Risk Management Using RAROC Models Christopher James, September 1996
96-47	Value at Risk: Implementing a Risk Measurement Standard Christopher Marshall and Michael Siegel
96-48	Evaluating Value-at-Risk Methodologies: Accuracy versus Computational Time Matthew Pritsker, November 1996
96-49	Calculating Value-at-Risk William Fallon, January 1996